

Google uses machine learning at data centers in efficiency drive

June 1 2014, by Nancy Owano



The mechanical plant at our facility in The Dalles, Ore. The data center team is constantly tracking the performance of the heat exchangers and other mechanical equipment pictured here.


How low can we go? That is a big question for a company such as Google which needs to think carefully about energy usage at data centers, the backbone of Google's entire business. Google aggressively pursues ways to carry out a more efficient use of power. Google, in fact, is not shy to talk about how proud it is toward its efforts. For over a

decade, said Google in a May 28 blog entry, they have been designing and building data centers that use half the energy of a typical data center, and they said they are always looking for ways to reduce energy use even further. The blogger, Joe Kava, vice president, data centers said in the company's pursuit of efficiency they have hit on a new tool, machine learning. Kava said the company has begun using a neural network to analyze the oceans of data it collects about its server farms and to recommend ways to improve them. Kava relayed on the blog that Google has developed a machine learning algorithm that learns from operational data to model plant performance and predict PUE.

What is PUE? The data-center industry uses power usage effectiveness to measure efficiency. A PUE of 2.0 means that for every watt of IT power, an additional watt is consumed to cool and distribute power to the IT equipment. A PUE closer to 1.0 means nearly all of the energy is used for computing. Google said their calculations include performance of their entire fleet of data centers worldwide. Their calculations occur year-round, not seasonally. In the end, they have a trailing twelve-month PUE of 1.12 across all their data centers, in all seasons, including all sources of overhead. (Google includes servers, storage, and networking equipment as IT equipment power and everything else as overhead power. Also, seasonal weather patterns affect PUE values, which are lower during cooler quarters. Google has kept a low PUE average across its data center sites even during hot, humid summers in Atlanta. That means the average PUE for all Google data centers is 1.12, even though they could have boasted a lower number if using narrower boundaries. Google also claimed the TTM energy-weighted average PUE for all Google [data centers](#) at 1.12, makes their centers among the most efficient in the world.)

On May 28, Google released a white paper, "Machine Learning Applications for Data Center Optimization" by Jim Gao, a Google engineer on the team. The paper discusses how neural networks are

being used to optimize data center operations and drive energy use to new lows. Kava emphasized praise for the contributions that Gao made, for having studied up on machine learning and then for having started building models to predict and improve data center performance.



A simplified version of what the models do: take a bunch of data, find the hidden interactions, then provide recommendations that optimize for energy efficiency.

Applying machine learning principles to data center monitoring was not trivial. Kava said what Gao designed works like other examples of machine learning, such as speech recognition, where a computer recognizes patterns out of large amounts of data and can "learn" from them. Nonetheless, in the data center environment, said Kava, "it can be difficult for humans to see how all of the variables—IT load, outside air temperature, etc.—interact with each other. One thing computers are good at is seeing the underlying story in the data, so Jim took the information we gather in the course of our daily operations and ran it through a model to help make sense of complex interactions that his

team—being mere mortals—may not otherwise have noticed."

In the paper, Gao said, "A [machine learning](#) approach leverages the plethora of existing sensor data to develop a mathematical model that understands the relationships between operational parameters and the holistic energy efficiency. This type of simulation allows operators to virtualize the DC for the purpose of identifying optimal plant configurations while reducing the uncertainty surrounding plant changes."

Kava said Gao's models are 99.6 percent accurate in predicting PUE. "This means he can use the models to come up with new ways to squeeze more efficiency out of our operations."

More information: * [www.google.com/about/datacente ... efficiency/internal/](http://www.google.com/about/datacente...efficiency/internal/)
* [googleblog.blogspot.com/2014/0 ... through-machine.html](http://googleblog.blogspot.com/2014/0...through-machine.html)
* [static.googleusercontent.com/m ... mization-finalv2.pdf](http://static.googleusercontent.com/m...mization-finalv2.pdf)

© 2014 Tech Xplore

Citation: Google uses machine learning at data centers in efficiency drive (2014, June 1)
retrieved 9 April 2024 from
<https://techxplore.com/news/2014-06-google-machine-centers-efficiency.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.
